Stefan Scholl, DC9ST

Malvenweg 3, 89160 Dornstadt, Germany: **dc9st@panoradio-sdr.de**

The Panoradio: A Modern Software Defined Radio with Direct Sampling

This SDR captures up to 100 MHz bandwidth at once and provides an extra-wide panorama view seamlessly from dc to VHF, as well as direct sampling of the 70-cm band.

Direct sampling receivers are considered as the ultimate concept of software defined radio (SDR) as they digitize the analog signal right after the antenna with a minimum of analog processing. This is a desirable approach, since digital signal processing has many advantages over analog signal processing.

In recent years several developments of direct sampling receivers occurred in Amateur Radio, such the "All-Digital Transceiver" from James C. Ahlstrom, N2ADR[1], the well-known WebSDR, from Pieter-Tjerk de Boer, PA3FWM[2] , the HPSDR project[3], SDR version of the Red Pitaya[4] of Pavel Demin, and many more including numerous commercial products. The sampling rates of many of these SDRs are usually in the range of 60 to 130 MHz and cover frequencies in the short wave band from 0 to 30 MHz or even up to 65 MHz. An overview with a selection of direct sampling receivers in Amateur Radio is shown in Table 1.

The direct sampling Panoradio SDR

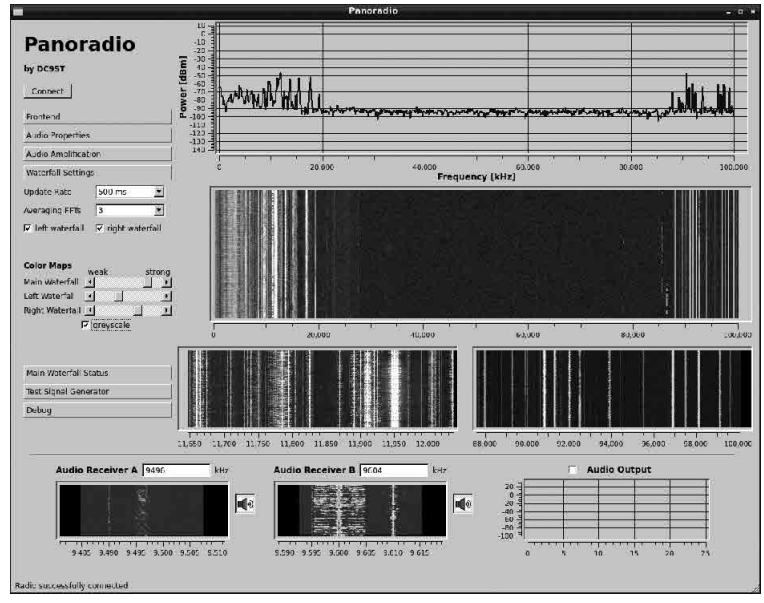

Figure 1 — Panoradio GUI shows the main waterfall in the full span from 0 to 100 MHz. Short wave signals occupy only a fraction at the left of the complete span. Signals above 88 MHz are FM broadcast stations and the sporadic emissions below 88 MHz belong to emergency services. The small left waterfall shows the 15 m broadcast band. The small right waterfall zooms on the VHF broadcasting signals.

Table 1.
Some examples of direct sampling receivers in Amateur Radio.

Direct Sampling SDR	*AD Converter*	*Sampling Rate (Msps)*	*Frequency Range (MHz)*
N2ADR All-Digital TRX (note 1)	ADS 5500 (14 Bit)	123	0 – 61
WebSDR (note 2)	LTC 2216 (16 Bit)	78	0 – 30
HPSDR (note 3)	LTC 2208 (16 Bit)	130	0 – 65
Red Pitaya SDR (note 4)	LTC 2145 (14 Bit)	125	0 – 50
Panoradio SDR	AD 9467 (16 Bit)	250	0 – 100, 425 – 445

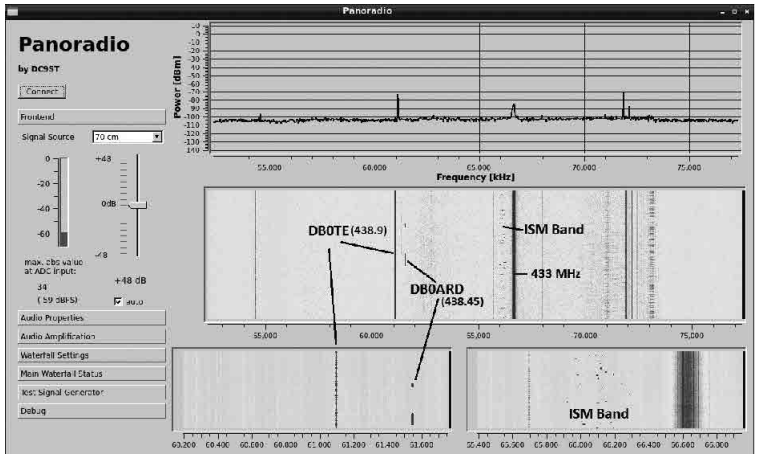

Figure 2 — Experimental reception of 70-cm band signals with undersampling shows various signals, such as the 433 MHz ISM band and two 70-cm band ham repeaters. Scales show the aliasing frequencies.

Figure 3 — Stand-alone operation with mouse and monitor connected directly to the SDR.

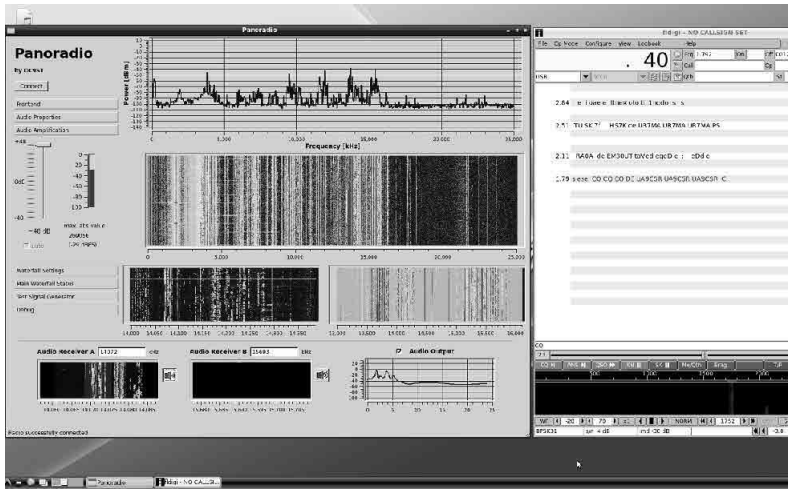

Figure 4 — Demodulation of PSK31 signals with Fldigi running directly on the SDR hardware.

receiver presented here digitizes the signal with a very fast 250 Msps sampling rate with 16 bits. It also has an experimental mode to receive UHF signals in the 70-cm band (425-445 MHz) using undersampling. The receiver has a unique panorama function. It captures and displays aliasing-free signals from 0 to 100 MHz in three zoomable waterfall displays for spectrum monitoring. These panoramic displays can independently zoom in to a bandwidth of 6 kHz and thus deliver a minimum resolution down to 7 Hz. They allow the observation of different bands in the complete frequency range from medium wave to VHF simultaneously. For example it is easy to watch signals in the 80 m, 10 m and 4 m band at the same time (Figure 1). A screenshot of the GUI during experimental reception of the 70-cm band can be seen in Figure 2.

In addition to the panoramic displays, two independent audio receivers tunable from 0 to 100 MHz are available. Each audio receiver covers a bandwidth of 22 kHz (with optional filters for 6 or 2.4 kHz), and displays the signal spectrum in additional small waterfalls. The two audio signals can be multiplexed or mixed prior to SSB demodulation using a Weaver demodulator.

The SDR is a standalone embedded system, where mouse, monitor and keyboard can be attached directly for full operation as shown in Figure 3. The hardware is built around the Zedboard, a modern embedded hardware platform. The Zedboard contains the Xilinx Zynq device, that combines a FPGA for fast data processing with a dual ARM A9 processor core running a Linux operating system. With the operating system running on the Zedboard, the software for controlling the radio and demodulation software, such as Fldigi[5], runs directly on the SDR (Figure 4). The inside view of the Panoradio is depicted in Figure 5.

The Panoradio is intended as a demo, that shows what is possible with today's technology in AD conversion and signal processing hardware. It is an open source project, so the design files can be accessed from the project website: **panoradio-sdr.de**.

The Panoradio consists of three parts: an analog front end, an AD converter and digital signal processing, which this paper presents in detail in the following sections.

Basics of Direct Sampling

First we briefly introduce and comment on direct sampling.

Direct Sampling Receivers

In a SDR receiver the AD converter marks the boundary between analog and digital signal processing (DSP). Historically the AD converter moved from the end of

the reception path (where bandwidth is low) further towards the front (where bandwidth is higher). Doing so, more and more analog circuits are replaced by DSP. This is desirable, because DSP is superior to analog processing. The ultimate concept is "direct sampling", sometimes called "all-digital receiver", or "digital radio". Then the antenna is more or less directly connected to the AD converter with no analog down conversion. Usually an analog (anti-aliasing) filter and sometimes an analog preamplifier are included.

Advantages of DSP

The extensive use of digital signal processing has many advantages over analog circuits: Analog processing is often limited by the laws of physics, that can hardly be overcome. Digital processing is limited only by circuit complexity — a better performance (sensitivity, dynamic range, spurs, agility, etc.) is achieved by more complex calculations and larger bit widths. Since semiconductor technology has continuously advanced following Moore's law, very complex systems can be built today and it is possible to achieve extraordinary accuracy and performance for digital signals with comparatively little effort.

The advantages of digital signal processing are apparent at different stages of the receiver:

— Digital FIR filters can be built with virtually any filter response. Low insertion loss, steep transition and high stop-band attenuation are only a matter of calculation resources. In many cases even medium-sized FPGAs can handle complex high-performance filters extremely well.

— Digital mixers and amplifiers are implemented as multipliers that do not introduce any spurs, such as harmonics and unwanted IMD and do not have any gain imperfection or other parasitic behavior.

— Digital oscillators based on direct digital synthesis (DDS) achieve extremely high spectral purity with low hardware complexity, and have virtually no spurs or harmonics. The frequency can be changed instantaneously without phase discontinuity.

— In general, effects of aging, impedance mismatch and all sorts of EMC issues (grounding, coupling between HF components, noisy power supply, etc.), that can bug analog designers, are completely eliminated in digital signal processing. The problem of additional noise that is introduced by most analog circuitry and measured as noise figure, is also widely eliminated in the digital domain. Noise is present only as quantization noise, caused by discrete sampling in amplitude. However, this quantization noise can be made arbitrarily low using larger bit widths.

— A sometimes overlooked property of digital radios is that they can be copied easily, since the radio usually consists of software or other design files, such as programming code, VHDL, FPGA design files or a GNU radio project on a PC. Before trying modifications, the designer can "backup" the radio and retrieve it, if the changes do not work.

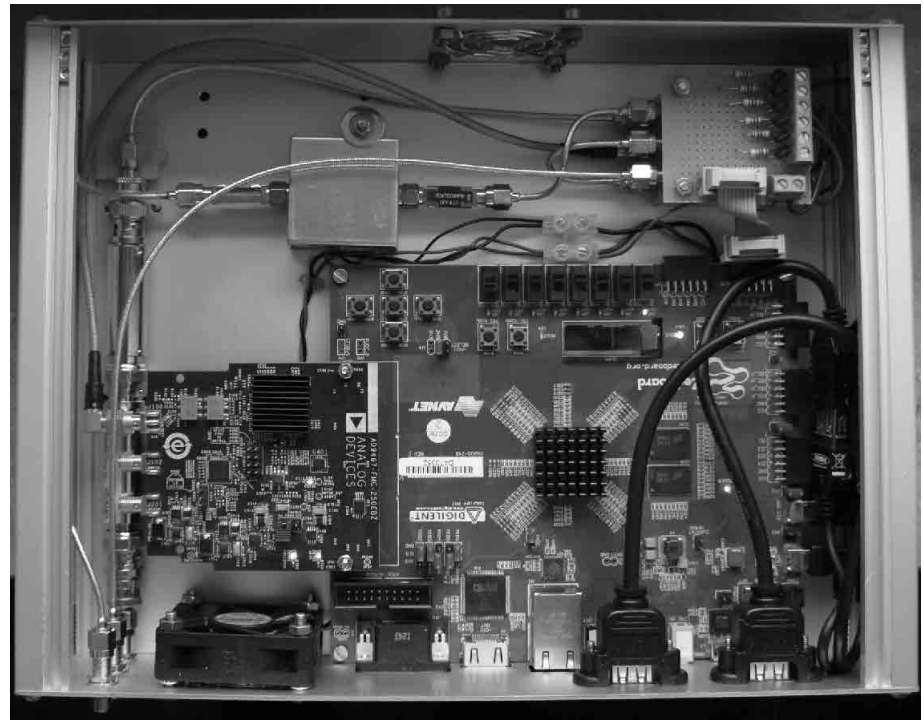

Figure 5 — The inside view shows the Zedboard (large center PCB) and the ADC board (small left-side PCB) with the anti-aliasing filter in the lengthy metal case on the left. The relays for switching the front end and their driver circuit are located In the upper right corner. The square metal box contains the 70-cm band front end.

Direct sampling receivers, indeed, have the potential to build high-performance receivers, since DSP can be implemented with almost arbitrary accuracy. Once the signal has been digitized, it is not prone to further degradation. The only bottleneck that may degrade the signal is the AD converter. It is important to carefully select the AD converter, since its properties define the overall performance of the radio in terms of linearity, sensitivity and bandwidth.

Challenges of Direct Sampling

When implementing a direct sampling receiver, the designer faces some special challenges regarding clock jitter and data rates.

First, the phase noise or clock jitter of the sampling clock is critical for high-speed SDRs. Even a small clock jitter may introduce additional noise in the ADC, which limits the SNR of the ADC. SNR is:

$$\mathrm{SNR} = 20\log\left(1/(2\pi f_a t_{jitter})\right) \qquad (1)$$

where t_{jitter} is the clock jitter in seconds and f_a is the maximum analog input frequency.

The SNR is not only dependent on the amount of clock jitter, but also increases with the frequency of the analog signal. Since direct sampling receivers often deal with higher analog frequencies than other SDRs, it is important to consider the noise introduced by clock jitter.

Second, the AD converter outputs the samples for further DSP at a very high data rate. The sampling theorem relates sampling rate with bandwidth: The sampling rate must be twice the highest frequency of the analog signal. The larger the processed analog bandwidth, the higher data rates are produced. In case of the Panoradio data rates up to 16-bit × 250 MHz = 4 Gbit/s occur at the input, and with I and Q, up to 2 × 22 bit × 250 MHz = 11 Gbit/s internally. This requires powerful hardware platforms like FPGAs because the computational load is too high for PCs or other processors.

Analog Front End and AD Converter

The Panoradio analog front end is mainly responsible for conditioning the analog signal coming from the antenna to prepare it for AD conversion, which includes proper anti-aliasing filtering.

Analog Front End

The Panoradio offers three analog input paths and analog front ends for different frequency ranges.

— A 90 MHz low-pass filter for a full panoramic view of HF and VHF.

— A 425-445 MHz band-pass filter plus a pre-amplifier for 70-cm band direct sampling reception.

— A direct path without filtering for experiments and flexibility.

AD converters in a direct sampling receiver require some analog front end to cope with aliasing due to sampling. Every analog input signal above $f_s/2$ appears in the band between 0 and $f_s/2$.

Wideband Reception: 0-100 MHz

To avoid aliasing in the Panoradio, with its $f_s = 250$ Msps, a low-pass filter must suppress all frequencies above 125 MHz before the AD converter. It is crucial to achieve a high stop-band attenuation, otherwise strong signals from TV transmitters or mobile phone services might interfere with signals in the 0 to 125 MHz range. Since most analog filters have a quite smooth transition from passband to stopband, the corner frequency of an anti-aliasing filter is usually significantly less than $f_s/2$, that is, less than 125 MHz.

In the Panoradio, two 90 MHz low-pass filters consisting of two cascaded Crystek CPFL 0090 filters are used. These filters provide around 55 dB stopband attenuation each, or 110 dB in total (Figure 6). With this high suppression, aliasing signals are completely eliminated before reaching the AD converter.

Undersampling for 70-cm band reception

The Panoradio has a second front end to enable undersampling of the 70-cm band. With undersampling the analog signal can have a frequency above $f_s/2 = 125$ MHz, which clearly violates the sampling theorem. So, aliasing occurs and these signals appear in the digital domain in the band below 125 MHz, according to the formula,

$$f_{alias} = \left| nf_s - f_{input} \right|$$

where f_{input} is the analog signal to be sampled and n is an integer, such that $f_{alias} < f_s/2$.

This behavior, which is not normally desired, provides an opportunity to translate signals above 125 MHz into the reception band, just as an analog mixer would do. The 70-cm band from 430 to 440 MHz, then translates to between 70 and 60 MHz, with an inverted spectrum.

Since in AD conversion many frequencies have the same alias frequency, as can be seen from the formula, analog signals potentially interfere with each other in the digital domain below 125 MHz. For example not only the band 430 – 440 MHz, but also the bands 180 – 190, 310 – 320, 560 – 570 MHz and so on, appears in 70 to 60 MHz. Therefore it is mandatory to use an analog band-pass filter before AD conversion that passes only signals from the 70-cm band, ensuring that no interference occurs.

The task of the Panoradio 70-cm band front end is twofold. First, it provides suppression of signals out of the 70-cm band. Second, it amplifies the input signal in order to improve sensitivity. The front end consists of a 433 MHz band-pass filter (Foxtech BPF 433) with 50 dB stop-band attenuation, followed by a low noise amplifier LNA4ALL[6] (+23 dB) and a second band-pass filter, shown in Figure 7. The first filter prevents the amplifier from overloading from unwanted out-of-band signals. The second filter sets the ultimate anti-aliasing and prevents broad band noise from the LNA to reach the AD converter, achieving a total amplification of +20 dB and more than 80 dB stopband attenuation (Figure 8).

High-Speed AD conversion

The AD converter is the heart of every direct-sampling receiver. Its properties are crucial for the overall receiver performance. Important ADC parameters are SNR, SFDR (spur free dynamic range), analog bandwidth and full-scale voltage. It should be noted that the number of bits is of minor importance, because the least significant bits usually carry only noise. The ENOB (effective number of bits) considers this issue and describes the number of bits that are "noise-free".

The Panoradio uses the AD9467-250, which is a state-of-the-art 16 bit, 250 Msps

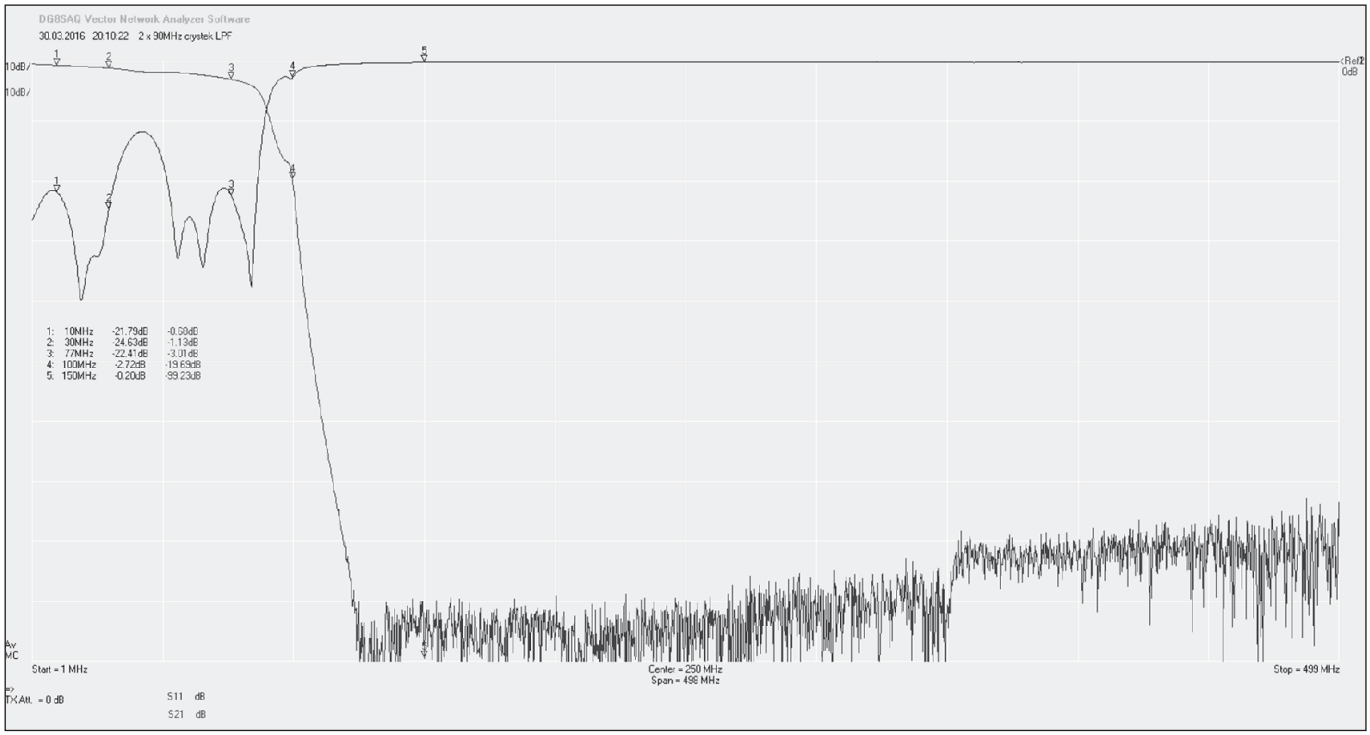

Figure 6 — 100 MHz front end with two cascaded analog low-pass filters. The dynamic range is limited by the VNWA 2 used for measurement.

AD converter with excellent properties and a wide analog bandwidth[7]. It is characterized by,

— SNR: 76 dBFS / 12.3 ENOB
— SFDR: 100 dBFS
— Analog bandwidth: 0 – 900 MHz
— Full scale voltage: 2.5 V p-p.

The Panoradio uses the AD9467 evaluation board from Analog Devices with a FMC connector, that perfectly fits into the Zedboard FMC socket. Minor soldering modifications were necessary on the evaluation board to enable the on-board clock source, and to route the analog signals correctly to the AD converter chip (circumventing the optionally available preamp and clock conditioner implemented on the board) and to provide impedance matching.

Impedance matching from a 50 Ω single-ended system to the ADC differential input with its 530 Ω in parallel with 3.5 pF consists of mainly a resistive matching plus some additional components for filtering of high-frequency noise. Resistive impedance matching has the advantage of providing a wideband match supporting frequencies up to many hundreds of megahertz. Two 1:1 transformers convert the signal from single-ended to differential. Using two transformers mitigates phase and amplitude imbalances that otherwise would cause spurs[8].

The full scale power of the sole AD converter at 2.5 V p-p for the 530 Ω input results in

$$P_{FS} = \frac{V_{p-p}^2}{8R} = \frac{2.5^2}{8 \cdot 530} = 0.00147 \text{ W}$$

which equals 1.7 dBm.

With the SNR of 76 dB and the thermal noise for a 125 MHz bandwidth of -93 dBm the ADC noise figure is around 19 dB.

However, taking the analog matching circuit of the ADC into account, the full scale voltage rises due to the resistive impedance matching network. Since a resistive match is not lossless, it introduces a loss of approximately 9 dB. Therefore the full-scale power of the complete circuit rises to $P_{FS} \approx$+11 dBm. Note that this decreases the sensitivity by 9 dB, resulting in a final noise figure of approximately 28 dB.

The full scale power of the AD converter circuitry of +11 dBm is quite large for the application of a direct sampling SDR. Noise figure is usually not a major problem for operation at low frequencies in the HF band. For higher frequencies it is expected to affect the receiver sensitivity. To cope with this shortcoming, that virtually all direct sampling receivers exhibit, a preamp is used to get an improved practical dynamic range and sensitivity.

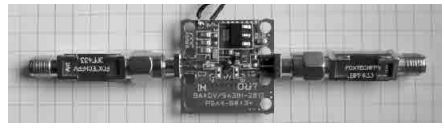

Figure 7 — 70-cm band front end with two 433 MHz band-pass filters and the LNA4ALL preamp.

For the Panoradio wideband input (0 – 100 MHz) I decided to not use a preamp to keep distortions low. An active preamplifier always introduces distortion, a problem that increases with the number of strong signals an ADC is exposed to. Due to the high bandwidth and large number of signals in the Panoradio, using no preamp excludes possible distortions. Transformers may also introduce distortions and considerable insertion loss at higher frequencies due to parasitic behavior. This effect gets worse for larger turn ratios, that is, higher voltage gain[8]. For ratios like 1:4 or 1:9 the maximum usable signal frequency is limited, which would heavily impair performance for 70-cm band reception. However, if high sensitivity is required, it is easy to attach an external preamp (preferably with band limitation filter). So the decision against a preamp or transformer in the Panoradio was in favor of wideband operation with as little non-linear distortion as possible.

Later FFT processing exhibits large processing gains such that the sensitivity for monitoring signals in a waterfall is much better than the specified ADC sensitivity (up to 73 dB for full zoom).

Low clock jitter is important to preserve the ADC SNR. The evaluation board features a high-performance Vectron VCC 6 clock generator that has a very low jitter of 133 *f*. Equation (1) shows the maximum achievable

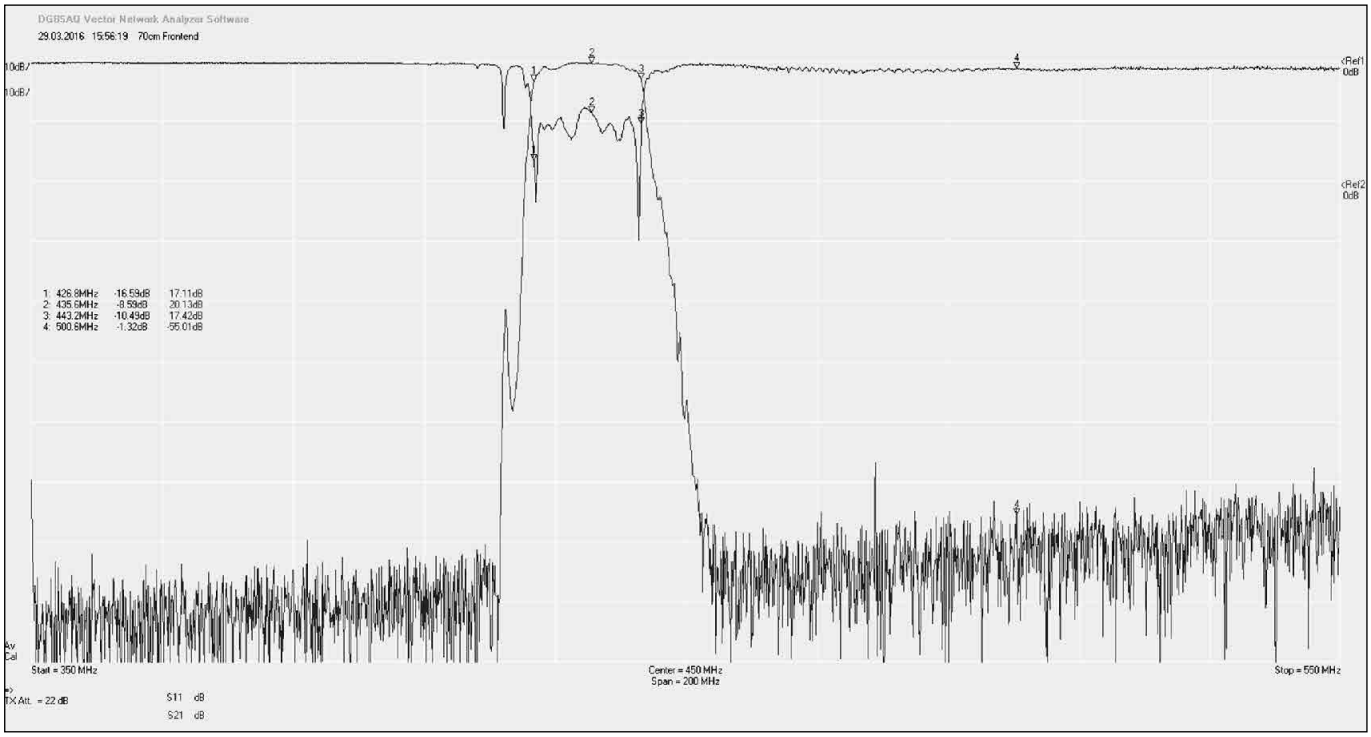

Figure 8 — 70 cm band front end measurements. The dynamic range limited by the VNWA2 used for measurement.

SNR due to clock jitter. For input frequencies from 0 - 100 MHz (the Panoradio main frequency range) the theoretical SNR is always more than 80 dB, which is clearly above the ADC SNR of 76 dB. So the clock jitter of the Vectron VCC 6 does not affect the noise performance of the receiver.

For input frequencies above 240 MHz, the impact of the clock jitter cannot be neglected. For 430 MHz (70-cm band), the ADC SNR and sensitivity is reduced by approximately 7 dB. Therefore, in contrast to the wideband input, the 70-cm band front end uses a preamp to make up for this additional loss by increasing the signal level and thus lowering the noise figure from an excessive 34 dB to moderate 13 dB. Non-linear distortion introduced by the preamp is not an issue, because it is a narrowband front end. The band-pass filters reject out-of-band signals before they reach the preamp. This also shows that sensitivity is a general issue for undersampling that needs to be countered by large amplification.

This completes the analog front end and the ADC discussion. From this point on the SDR does all processing with DSP.

The Zedboard as Digital Processing Platform

SDR algorithms run on a hardware platform, either as software on a processor such as a PC or microprocessor, or as a digital circuit in a microchip, like an FPGA. These different hardware platforms have diverse processing speed, flexibility and ease of programming. Direct sampling receivers usually deal with large bandwidths and high data rates, and therefore require powerful processing platforms. Before describing the Panoradio DSP implementation details, the available hardware platforms and requirements are further analyzed.

Processor vs. FPGA

A processor is basically a digital circuit that executes software serially and is therefore inherently slow and not suitable for large bandwidth processing in high-end SDRs. Exceptions are multi-core processors, like GPUs in graphics cards, that exhibit parallelism to some extent.

The advantage of the processor is its flexibility, since the tasks it fulfils can be changed easily by executing different software. Also, it is comparatively easy to program software. It is possible to run an operating system with the advantage of easy access to user interfaces like a monitor, keyboard, mouse, and providing access to hard disks for storage or for an internet connection.

In SDR applications, processors are very good for:

— demodulation of different modes

— providing a graphical user interface (GUI)

— controlling the receiver functionalities

— recording and storing signals

— providing network and internet access.

FPGAs are fundamentally different from processors. An FPGA is a microchip that contains a large number of configurable circuit elements, such as adders, multipliers, multiplexers and registers, that can be used to create complex digital circuits. For every algorithm, that is to be executed on the FPGA, a specialized digital circuit must be designed and loaded onto the FPGA.

The big advantage of FPGAs is that they can process highly parallel data, and can easily deal with data rates of several hundred Msps. The achievable speed-up over processors is often in the order 100 to1,000 times, even though FPGA clocks are often much slower than processor clocks. Since the circuit in an FPGA is custom designed, the bit widths of data, memories and arithmetic units can be application specific, which is very efficient. If, for example, a bit width of 6 bits is considered sufficient for the data, the processing units are laid out to use only 6 bits, whereas in a processor the bit width is restricted to a predefined number such as 24, 32 or 64 bits.

The drawback of FPGAs is that the hardware implementation (circuit design) of algorithms requires good knowledge of hardware description languages, digital circuit design and experience with the design tools. Also the implementation of interfaces to the outside (monitor, keyboard, network connection) is difficult.

In SDR applications, FPGAs are very good for high-speed, repeatedly occurring calculations like, mixing and down conversion; filtering; high-speed FFT; and parallel demodulation of a large number of signals.

In summary, processors and FPGAs each have advantages and disadvantages. In a SDR a wide range of different tasks must be performed with different speeds. That is difficult to achieve with only a processor or only an FPGA. So the ideal hardware for SDR is a combination of both — and indeed, high-end SDRs already follow this configuration. So also does the Panoradio, which uses the "Zynq" device from Xilinx, that combines a processor and an FPGA on a single chip.

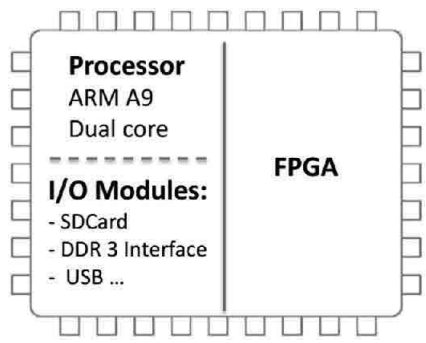

Figure 9 — The Zynq SoC features a processor and an FPGA as well as I/O interfaces on a single chip.

The Zynq

The Zynq chip is the result of a recent development in microelectronics to combine processors and FPGAs for acceleration in a system on chip. It contains a dual ARM Core A9, an FPGA and various interfaces (such as DDR3 memory interface, USB, SD Card) as shown in Figure 9. There are several thousand possible internal connections between the processor and the FPGA, which are programmable and allow fast transmission of data. This is a big advantage over a separate combination of FPGA and processor on a PCB, where the number of possible connections is small and can not easily be changed after PCB design.

The interconnect between processor and FPGA is a memory mapped interface. That means that a part of the processor memory is virtually connected to a 32-bit register in the FPGA through a AXI bus (Figure 10). The processor can write data to these special memory addresses (using the "mmap" command or a Linux driver), that is then transferred to the FPGA registers. Or it can read the registers in the FPGA by reading data from the memory. Direct memory access (DMA) is available for transferring large amounts of data.

Programming the Zynq involves two tasks; programming the processor, and designing the digital circuit for the FPGA. These two tasks can be done nearly independently by different tools.

The processor can be programmed either bare metal (programming it directly in *C*) or using an operating system. I strongly recommend using an operating system like Linux, which hides many details of the processor and provides easy access to user interfaces as well as to standard software. Then any programming language can be used to develop software for the Zynq.

The hardware design for the FPGA is done with Xilinx software "Vivado" using VHDL or Verilog and predefined IP blocks. For the Panoradio, the different receiver blocks (like FIR filters and DDS oscillators) are implemented using VHDL plus Xilinx IP cores. These blocks are connected together in a top level block-design schematics by using the graphical editor provided in Vivado.

The Zedboard

The Zedboard is a low-cost and open-source evaluation board and is ideal for experimenting with the Zynq device (**www.zedboard.com**). The Zedboard contains many peripherals that unleash the power of the Zynq, such as power supplies, 512 MB DDR3 Memory, Gigabit Ethernet, USB connectors, HDMI interface, SD card slot and an audio interface (sound codec ADAU1761). These peripherals make the Zedboard a fully stand-alone system. Additionally, the Zedboard offers a FPGA Mezzanine Card (FMC) connector, which is a standardized socket for attaching peripheral boards. I used this socket to attach the ADC board.

Receiver DSP Implementation

This section focuses on what is actually implemented in the Zynq FPGA and processor. SDR signal processing includes many different steps: digital down conversion, filtering, demodulation, FFT, control, interfaces and GUI. High-speed data processing is handled in the FPGA part of the Zynq, whereas low data rate tasks are handled by the processor. The overall implementation is shown in Figure 11.

After AD conversion the FPGA receives the samples via the ADC data interface module. It performs some ADC tests, takes care of line delays and converts the differential double-data-rate signals to single-ended data with single data rate, that can be processed in the FPGA. Optionally, a test signal generator can feed well-defined waveforms into the circuit for testing. Three digital down conversion (DDC) blocks perform the actual reception. One DDC implements functionality for the three zoomable waterfall displays. Two DDCs are dedicated to the two audio receivers. The following audio post processing block can further reduce bandwidth and contains a Weaver SSB demodulator. Interfaces for audio and HDMI to the processor are also implemented in the FPGA.

The ARM A9 processor takes care of

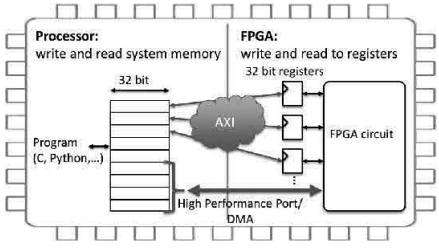

Figure 10 — Data transfer between the processor and the FPGA via memory mapped AXI bus or DMA.

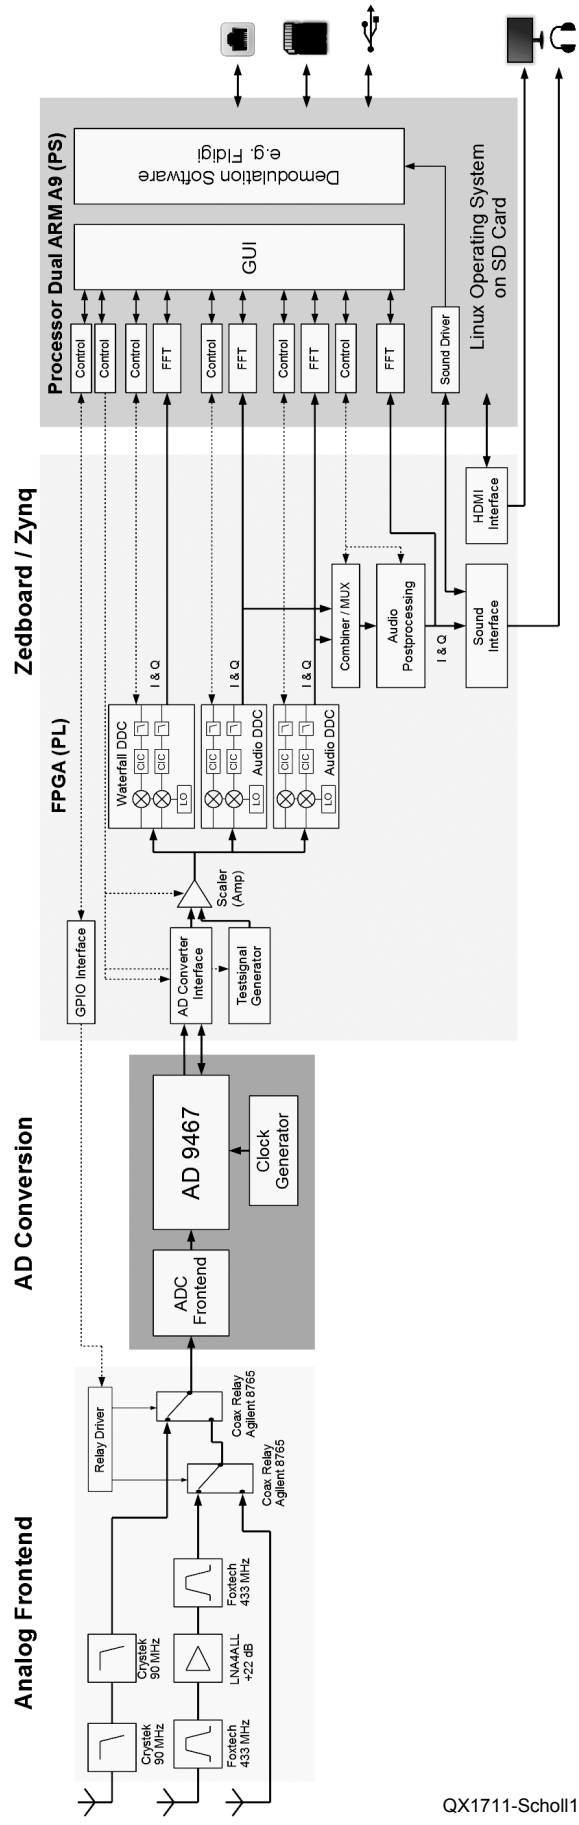

Figure 11 — Block diagram of the complete receiver showing the three main parts: analog front end, AD converter, and the Zedboard for digital processing.

all slow-speed processing, such as running a GUI to control the radio, calculating FFT, displaying received signals and executing demodulation software (Figure 4). Moreover it runs the OS for standalone operation.

DDC for Zoomable FFT

The three Panoradio waterfall displays can independently zoom on signals from the full span of 100 MHz to a span of 6.1 kHz as is depicted in Figure 12. This task is accomplished by the implementation of a zoom FFT with flexible bandwidth. This is a DDC, that down converts and filters the selected portion of the spectrum followed by an FFT. For this purpose, two challenges have to be solved: how to implement the zoom function efficiently, and how to implement three waterfalls efficiently.

DDC and Zoom Functionality

Overview: The DDC for the zoomable FFT is shown in Figure 13. It follows the typical design of a DDC with a digital local oscillator (LO), a complex mixing stage with two multipliers and CIC and FIR filters in the I and Q paths. The filters reduce the bandwidth to the desired span of the current waterfall setting. Finally, the I and Q samples are stored in a memory and sent to the processor for FFT.

Local Oscillator (LO): The implemented LO generates sine and cosine waveforms using direct digital synthesis (DDS). In DDS a sine is generated using a large look-up table. It stores samples of a sine waveform, which is then samples and interpolates at the desired frequency. A second table or a second address pointer generates the corresponding cosine. DDS LOs can achieve high accuracy and agility. For FPGA implementation I used the predefined Xilinx DDS IP block[9].

Filtering: CIC filters are low-complexity circuits that combine filtering with decimation. Decimation reduces the sampling rate from 250 Msps to a fraction determined by the decimation factor between 4 and 4,096 in order to relax the computational load for the FPGA in the subsequent stages. The following sharp FIR filters with 140 coefficients determine the final bandwidth and an additional decimation by a factor 8 sets the final sampling rate. The filter design has been done with the Matlab "Filter Design and Analysis Tool". The calculated filter parameters and coefficients were imported into the Xilinx IP blocks for CIC and FIR filters in Vivado.[10]

Zoom functionality: The required bandwidth of the DDC needs to be very flexible between 100 MHz and 6 kHz to implement the waterfall zoom function. For the cases of 100 MHz (1× zoom) and 50 MHz (2× zoom) bandwidth, the input samples are directly routed to the memory for a real FFT. For higher zoom levels, greater than 4× (span < 25 MHz), the signal passes through the DDC. The flexibility in output bandwidths from 25 MHz (4× zoom) to 6.1 kHz (16,384× zoom) is achieved only by a programmable CIC decimation rate of 4 to 4096 (16× and higher) or bypassing the CIC filters (4× and 8× zoom). By changing only the CIC decimation rate by a factor of 2, the output bandwidth also changes by a factor of 2. Thus the FIR filter remains fixed, which greatly simplifies the DDC. Figure 14 shows the resulting filter responses for different zoom levels.

Performance: The DDC is designed for excellent performance. The DDS provides LO signals with 21 bits each and 110 dB SFDR. The stop-band attenuation of CIC and FIR filters is typically greater than 105 dB (see Figure 14). The 140 coefficients for FIR filtering have a large bit width of 24 bits, such that the influence of coefficient

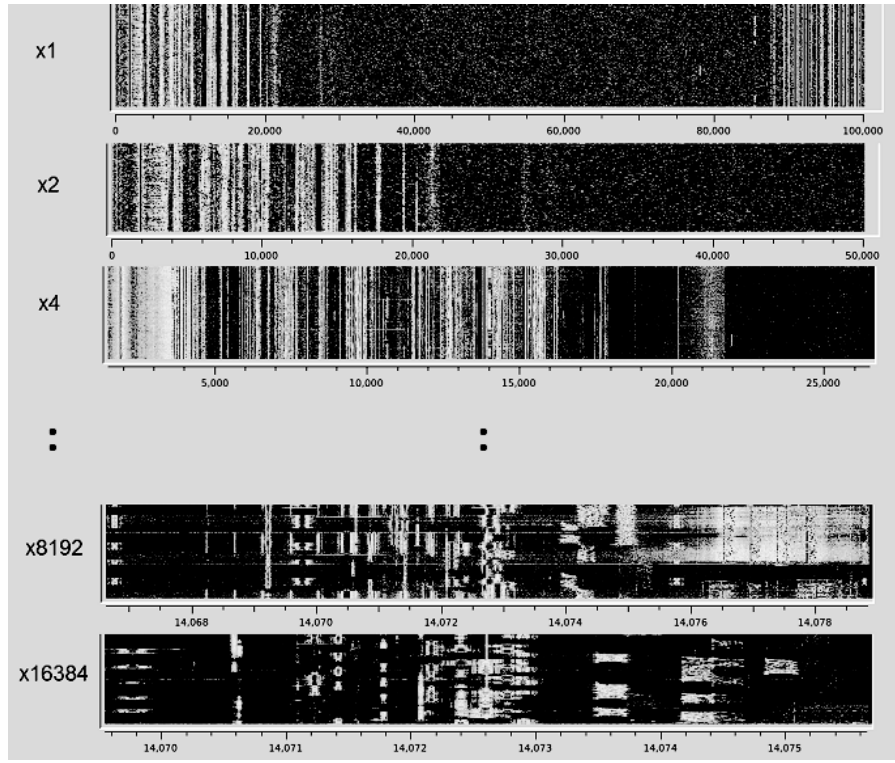

Figure 12 — Zoom function of the waterfall, here from full span to digital signals in the 20 m amateur band.

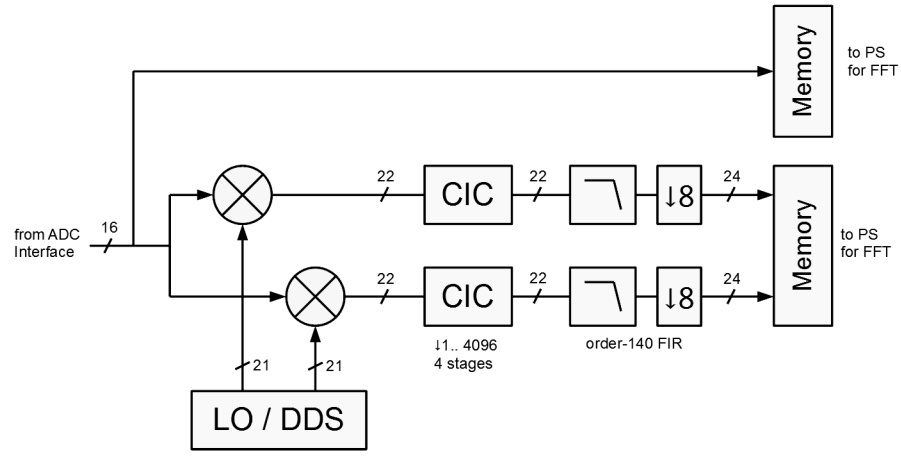

Figure 13 — DDC for zoomable FFT spectrum display.

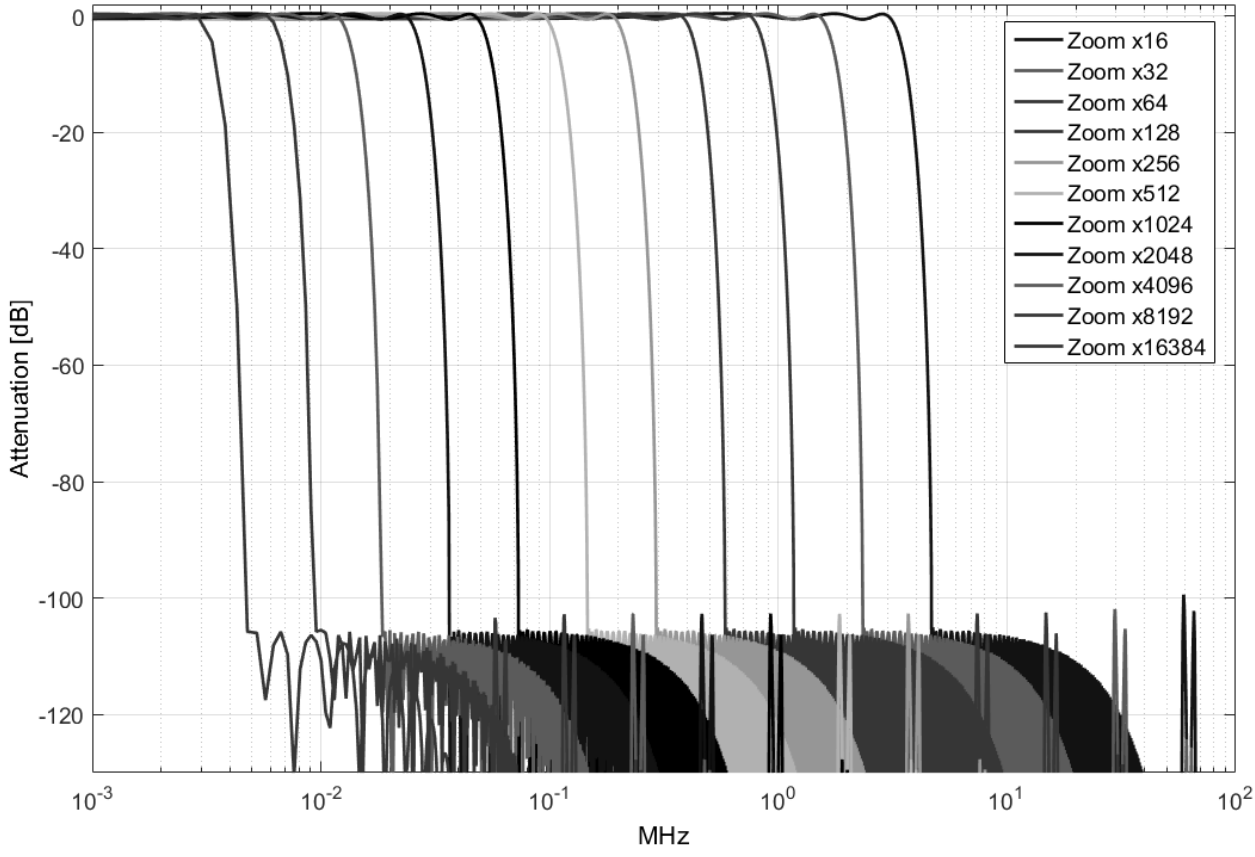

Figure 14 — Response of the DDC for the zoomable FFT for different zoom levels from x16 to x16384 achieved by only changing the decimation in the CIC from 4 to 4096.

quantization is negligible. For the overall DDC even worst case, for example full scale, out-of-band signals appear only with very small amplitude with less than 105 dBFS. In more realistic cases the DDC will not introduce any spurious signals. The bit width of I and Q signals is large enough to keep the influence of quantization errors very low (the introduced quantization noise is less than few tenths of a decibel). Note, that the bit width of I and Q channels increase with smaller bandwidth because noise power reduces, and therefore SNR increases.

One DDC – Three Waterfalls

For the three waterfall displays, three DDCs would be required. However, it is more efficient to share one single DDC between the three waterfalls. Then the DDC quickly switches continuously between the parameters of the different waterfall displays, captures data and calculates an FFT before switching to the next. The parameters that need to be switched are the center frequencies (LO) and CIC decimation rate. The LO frequency can change instantaneously since the DDS is very agile. The filters however require some time to recover from transients, which is a potential problem especially for high decimation values. However, worst case simulations of transient behavior showed, that the system recovers very quickly and this is not an issue if the switching time stays above several tens of milliseconds.

FFT Processing

For calculating the FFT, the DDC captures 8,192 I and Q samples in a memory (4,096 samples for 100 and 50 MHz bandwidth). Then the samples are weighted with a Hann window to reduce FFT leakage. A complex FFT is carried out resulting in 8,192 frequency bins. Since the monitor resolution is too low to display a 8,192 point spectrum, the spectrum is scaled down to 1,024 points by summing eight neighboring bins. This approach has a big advantage. It largely increases the resolution and sensitivity for detecting signals that are otherwise masked by the "leakage bell" of strong signals, see Figure 15. Improvements of some ten dB might occur in the vicinity of strong signals.

Optionally, the Panoradio can do several data captures and FFT operations consecutively to display the average spectrum. This does not reduce the noise floor, but reduces fluctuations in the spectrum plot.

Audio Reception and Demodulation – Audio DDC

The audio DDC follows the basic principle of the waterfall DDC with the exception, that the CIC has a fixed decimation rate of 2,048, see Figure 16. A FIR filter reduces the bandwidth to 11 kHz, so the total bandwidth for I and Q is 22 kHz, which fits well to a 48 kHz rate sound card. The audio I and Q outputs are passed to the audio post processing block for further bandwidth reduction and SSB demodulation. In addition a memory is included to capture samples for a 256 point FFT for the audio receiver waterfall displays. The additional decimation-by-two before FFT memory storage reduces the sample rate and thus increases FFT resolution. This decimation is not applied to the I and Q outputs, so as not to violate the sampling theorem at the subsequent SSB demodulator output. The audio DDC suppresses out-of-band signals by more than 80 dB, which provides very good attenuation of unwanted signals for practical operation.

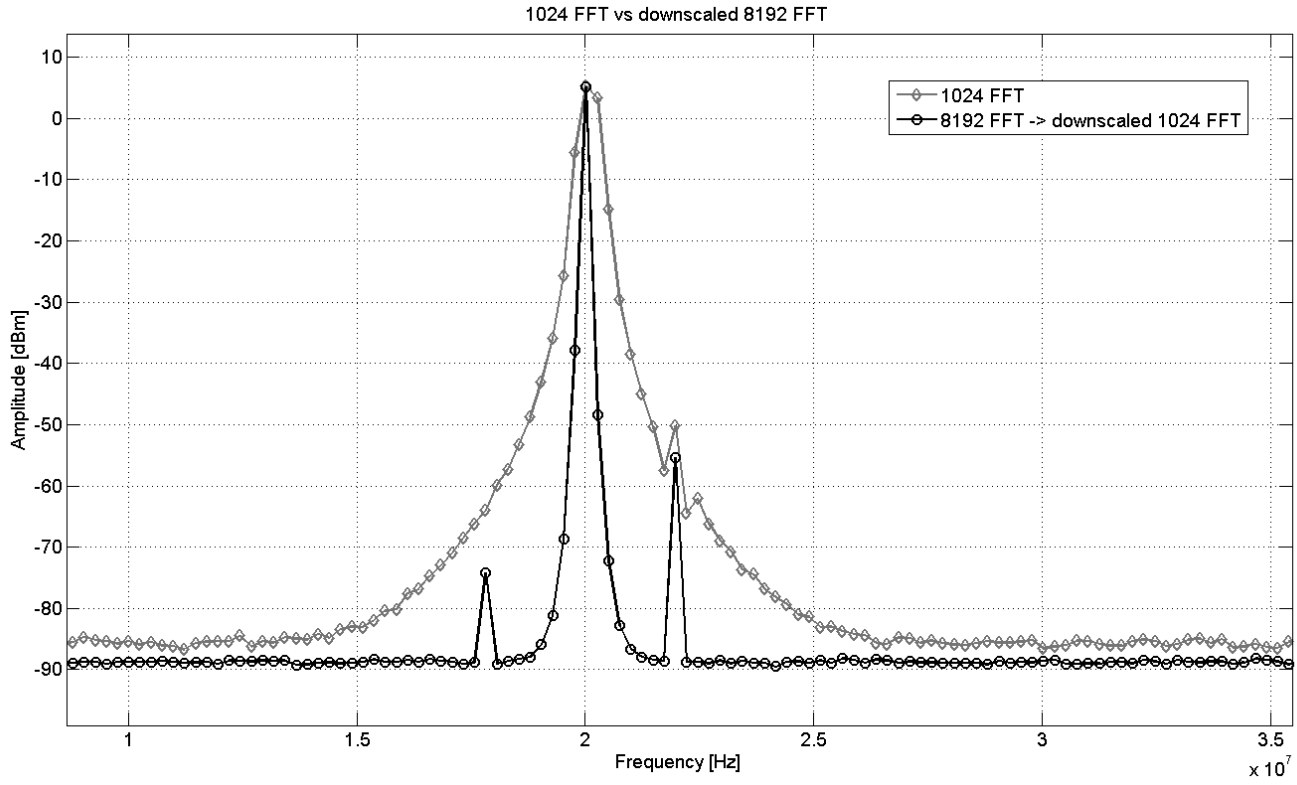

Figure 15 — Simulation and comparison of a native 1024-point FFT and a 1024-point FFT that originates from a downscaled 8192-point FFT.

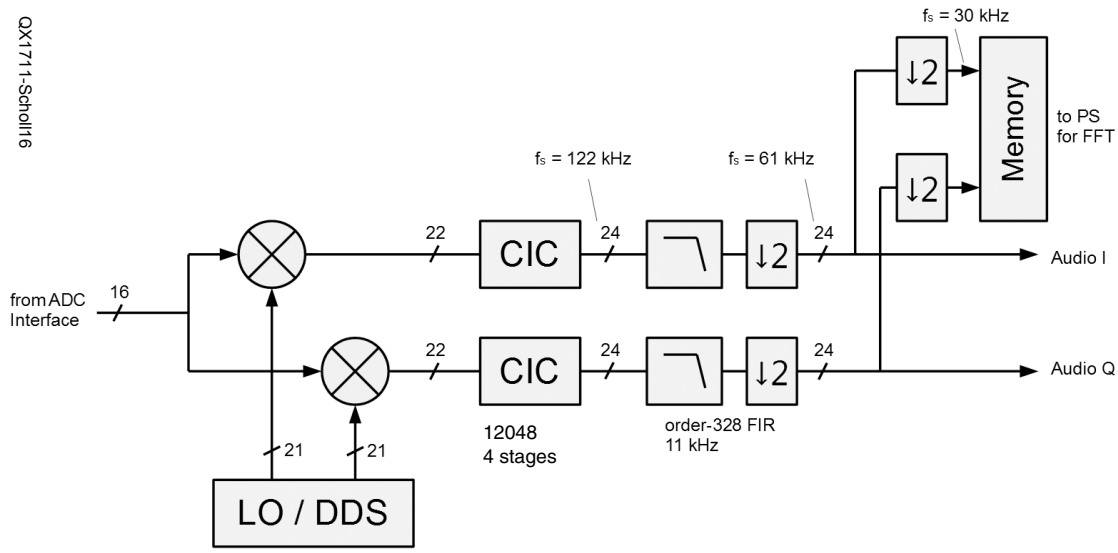

Figure 16 — One of the two audio DDCs implemented for audio reception.

The Panoradio features two of the above described audio DDCs. The I and Q output of the two DDCs are either multiplexed or summed. Summing allows for audio reception of two different bands simultaneously, that is, to receive and demodulate PSK31 signals from two bands at the same time.

Audio Post Processing and SSB Demodulation:

The I and Q samples from the audio DDC are fed to the audio post processing for volume setting, bandwidth reduction, SSB demodulation and sample rate conversion as shown in Figure 17. The first step is a clock domain crossing reducing the FPGA clock frequency from 250 MHz to 100 MHz, which greatly relaxes timing constraints for FPGA synthesis and place and route. This is possible because the signal bandwidth at this point is just several tens of kHz wide and the sampling rate is just 61 kHz. The clock domain crossing is followed by a software controlled variable gain amplifier for audio volume setting. Additional optional FIR filters can reduce the ultimate audio bandwidth from 22 kHz to 6 or 2.4 kHz. The signals can then be passed as I and Q values to the sound card directly or with prior SSB demodulation. A Weaver demodulator[11] is implemented for optional SSB demodulation. I consider the Weaver demodulator as one of the best options for SSB demodulation in DSP, since the required mixers, oscillators and low-pass filters are of low complexity and can be implemented with virtually perfect phase and amplitude balancing.

The sample rate after filtering (or optional SSB demodulation) is 30 kHz or 60 kHz depending on the selected filters. In any case this does not match with the 48 kHz sample rate of the Zedboard audio codec. Therefore an additional asynchronous sample rate converter interpolates the input signal at 30 or 60 kHz and filters the output to retrieve the signal at the desired 48 kHz.

FPGA Design

The high-speed digital processing parts of the Panoradio are implemented on the FPGA. The Zedboard Zynq 7020 is a mid-sized Zynq SoCs with 53,200 look-up tables and 106,400 flip-flops, 560 KB SRAM and 220 DSP slices. The FPGA for the Panoradio design is less than 50% occupied (46% look-up tables, 34% flip-flops, 39% memories, 45% DSP slices / multipliers). This quite low utilization would allow for an even more complex DSP. Main contributors to the FPGA utilization are the waterfall DDC, which operates at high sample frequencies and requires high parallelism (especially DSP slices for multiplications for FIR filtering), and the interconnect circuitry between processor and FPGA.

Control Software & Linux OS

The dual core ARM A9 processor runs the Linux operating system Linaro (based on Ubuntu for embedded devices). The OS and file system are stored on an SD card. With the Linux OS the Panoradio offers comparable possibilities as a PC (although the ARM A9 does not have as much computing power). The SDR can run any general purpose software for Linux. For this project, I compiled Fldigi for the ARM A9. Other useful software, like file browser, editors, audio players or the program "Scilab" for scientific calculations can run on the processor. I used QtCreator as the programming IDE. It offers the possibility to program directly on the Zynq instead of applying a cross compiler. Furthermore, the OS provides Ethernet, USB and HDMI interfaces for attaching network, monitor and input devices. During system start-up, a bootloader programs the FPGA with the bit stream containing the FPGA implementation.

The task of the software for controlling the Panoradio is mainly to provide a GUI and to control the radio functionality. This control includes setting the analog front end path, initializing the ADC and audio interface, setting the DDC frequencies and filter properties, determining amplifications for ADC and audio interface, and reading I and Q data from the DDCs for the FFT. Furthermore, the

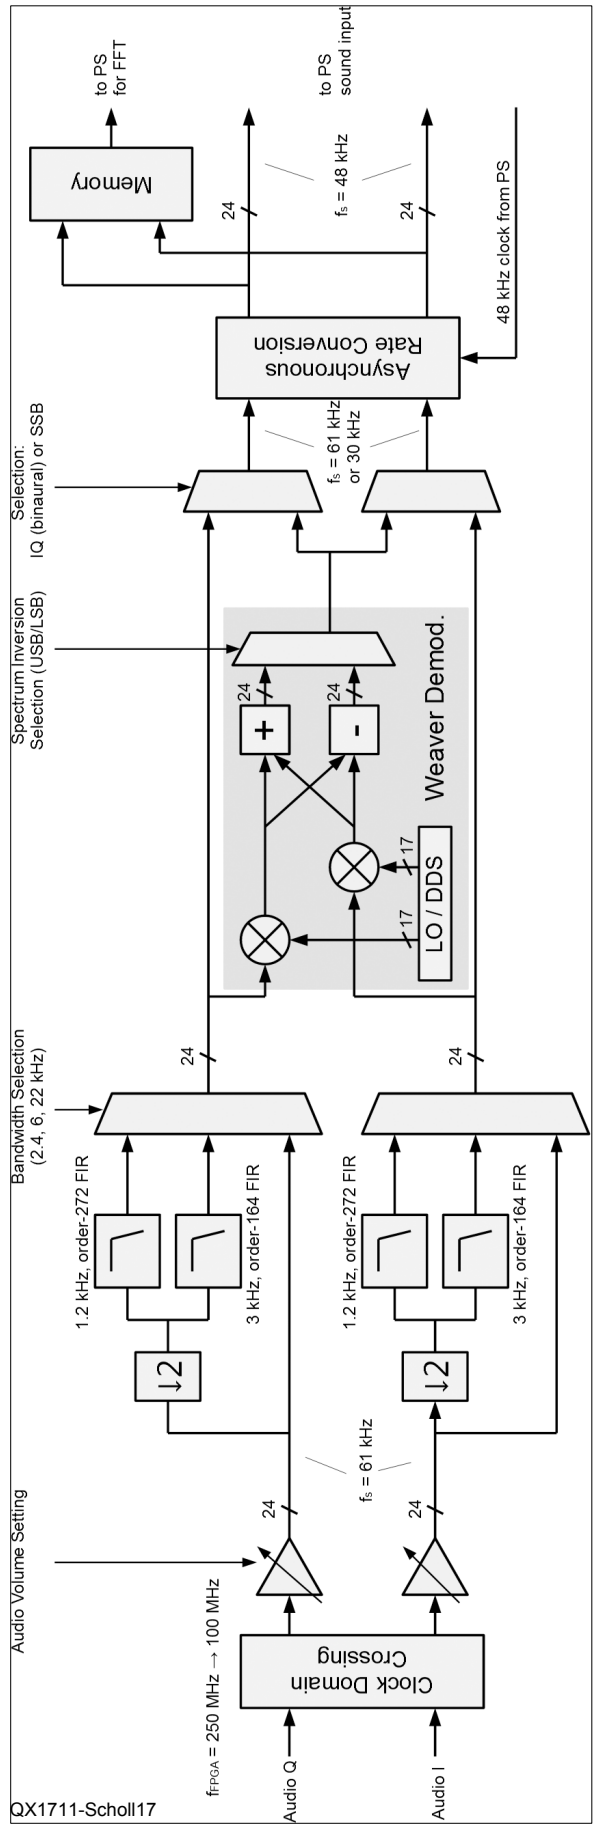

Figure 17 — Audio post processing: volume setting, final bandwidth selection, SSB Weaver demodulation, and sample rate conversion.

FFT calculations (including windowing, averaging, decimation, etc.) for the waterfall displays are done in software with the processor.

The control software is written in C++ and Qt. Qt provides a powerful framework for implementing a GUI and enhances C++ in many different useful ways. Qwt[12], a library for scientific applications, that provides various kinds of plots, is used to implement the waterfall displays.

The software GUI has the following features.

— Selection of the analog front end (100 MHz, 70-cm band or direct).

— AD converter input level monitoring.

— AGC control after the ADC.

— FFT averaging, update, time, and color scheme.

— Different modes for audio reception (bandwidth, SSB, spectrum inversion).

— Monitoring of audio signal level and setting for audio amplifier.

— Selection of a single audio receiver or the sum of both.

— Control of digital test signals.

— Monitoring of internal receiver states for debugging (signal and its spectrum at the different stages, setting of filters, LO, etc.).

For FFT calculation the software implements a wrapper for the FFTW library[13] a C library with fast FFT algorithms. The wrapper adds functions for applying a window, FFT averaging, converting to power scale in dBm and provides C++ style access to the FFTW library.

I chose C++ as the programming language, because the execution speed of the software is critical. Interestingly, the bottleneck is not the FFTs or the communication with the IP cores, but the drawing routines for the waterfall plots. The Zynq does not have any graphics acceleration core, which could speed up the drawing process. However, Xilinx has realized this bottleneck and included graphics acceleration in the Zynq successors; the UltraScale MPSoc, and the Zynq UltraScale+.

Conclusion

The Panoradio was designed to demonstrate the potential of today's analog and digital devices for SDR. Modern AD converters can digitize signals both fast and accurately. Recent digital FPGA devices can process the high data rates of direct sampling receivers and provide the possibility to run software and even a Linux OS on a single chip. The combination led to a stand-alone SDR covering frequencies from dc to VHF as well as 70-cm UHF band with direct sampling. The Panoradio is an open-source project available on **www.panoradio-sdr.de** in order to encourage Amateur Radio enthusiasts to try, learn and further develop direct sampling SDR technology.

Stefan Scholl, DC9ST, was originally licensed in 2006, as a German Class A license. He received the Diploma in electrical engineering in 2010 from University of Kaiserslautern, Germany, where he worked for 6 years as a researcher in the area of forward error correction for communication systems. He has nearly completed his PhD. Currently he is an engineer for radar technology at Hensoldt (formerly Airbus Defence and Space). Stefan's interests include homebrew equipment and hardware and algorithms for software defined radio.

Notes

[1]James C. Ahlstrom, N2ADR, "An All-Digital Transceiver for HF", *QEX* Nov./Dec. 2010.
[2]Pieter-Tjerk de Boer, PA3FWM, WebSDR: **https://wwwhome.ewi.utwente.nl/~ptdeboer/ham/sdr/**.
[3]High Performance Software Defined Radio, **www.openhpsdr.org**.
[4]Pavel Demin, Red Pitaya SDR, **pavel-demin.github.io/red-pitaya-notes/sdr-transceiver/**.
[5] Fldigi by W1HKJ, **https://sourceforge.net/projects/fldigi/files/**.
[6]Adam Alicajic, 9A4QV, **lna4all.blogspot.de/**.
[7]Analog Devices, AD 9467 Datasheet, Rev D ed.
[8]Rob Reeder (Analog Devices), "Transformer-Coupled Front-End for Wideband A/D Converters", Analog Dialogue, Vol. 39, Apr. 2005.
[9]Xilinx, DDS Compiler v6.0 LogiCORE IP Product Guide, PG141 ed., Jun. 2015.
[10]Xilinx, FIR Compiler v7.2 LogiCORE IP Product Guide, PG149 ed., Jun. 2015.
[11]Donald K. Weaver, "A Third Method of Generation and Detection of Single-Sideband Signals", Proc. of JRE, Nov. 1956.
[12]U. Rathmann. (2015, August) Qwt - Qt Widgets for Technical Applications. **qwt.sourceforge.net/**.
[13]M. Frigo and S. Johnson, "FFTW," MIT, Tech. Rep. 3.3.3, Nov., 2012.

used - surplus - obsolete - buy - sell - swap - all for free

no fees ~ no commissions ~ it's free !

www.SecondHandRadio.com

if it's electronic or electrical - find it here or sell it here

Announce your:
- For Sale
- Wanted
- Service
- Hamfest
- Swapmeet
- Contest
- Club Meeting
- Special Event
- Net skeds

Audio
Video
Tape
Phonograph
Wire
Connectors
Vintage
Tubes
Switches & Knobs
Books
Schematics
Manuals

Transmitters
Receivers
VFOs
Test Equipment
Hardware
Telephone
Tools
Inductors
Capacitors
Antiques
Solar
Wind

- Browse by state / province
- Remove your ad anytime
- Update your ad anytime
- Choose auto-expiration
- Link back to your website
- Dealers are welcome !
- Simple search
- Advanced search
- No limit on number of ads
- Upload 40 photos per ad

*see website for policies and restrictions

www.SecondHandRadio.com Since 2003

NOT an auction site; you are in complete control of your sale.